



Argonne
NATIONAL
LABORATORY

... for a brighter future



www.ultravis.org



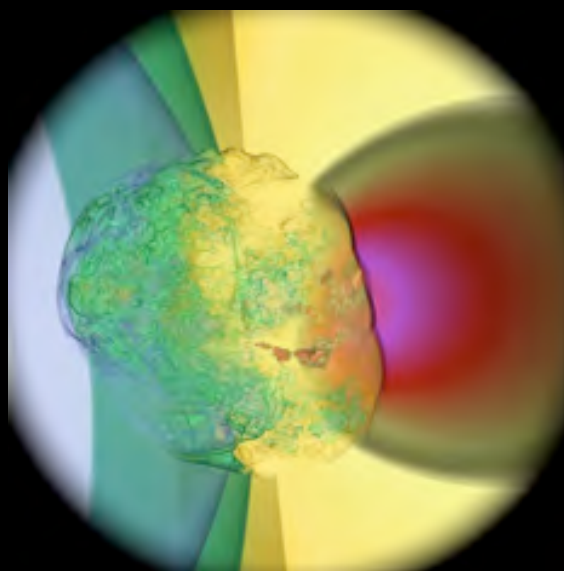
U.S. Department
of Energy

UChicago ►
Argonne_{LLC}



A U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC

Moving Analysis to the Data: Scalable Visualization Using Simulation Resources



Volume rendering of x-velocity in time-step 1530
of a hydrodynamics simulation of a core-collapse
supernova.

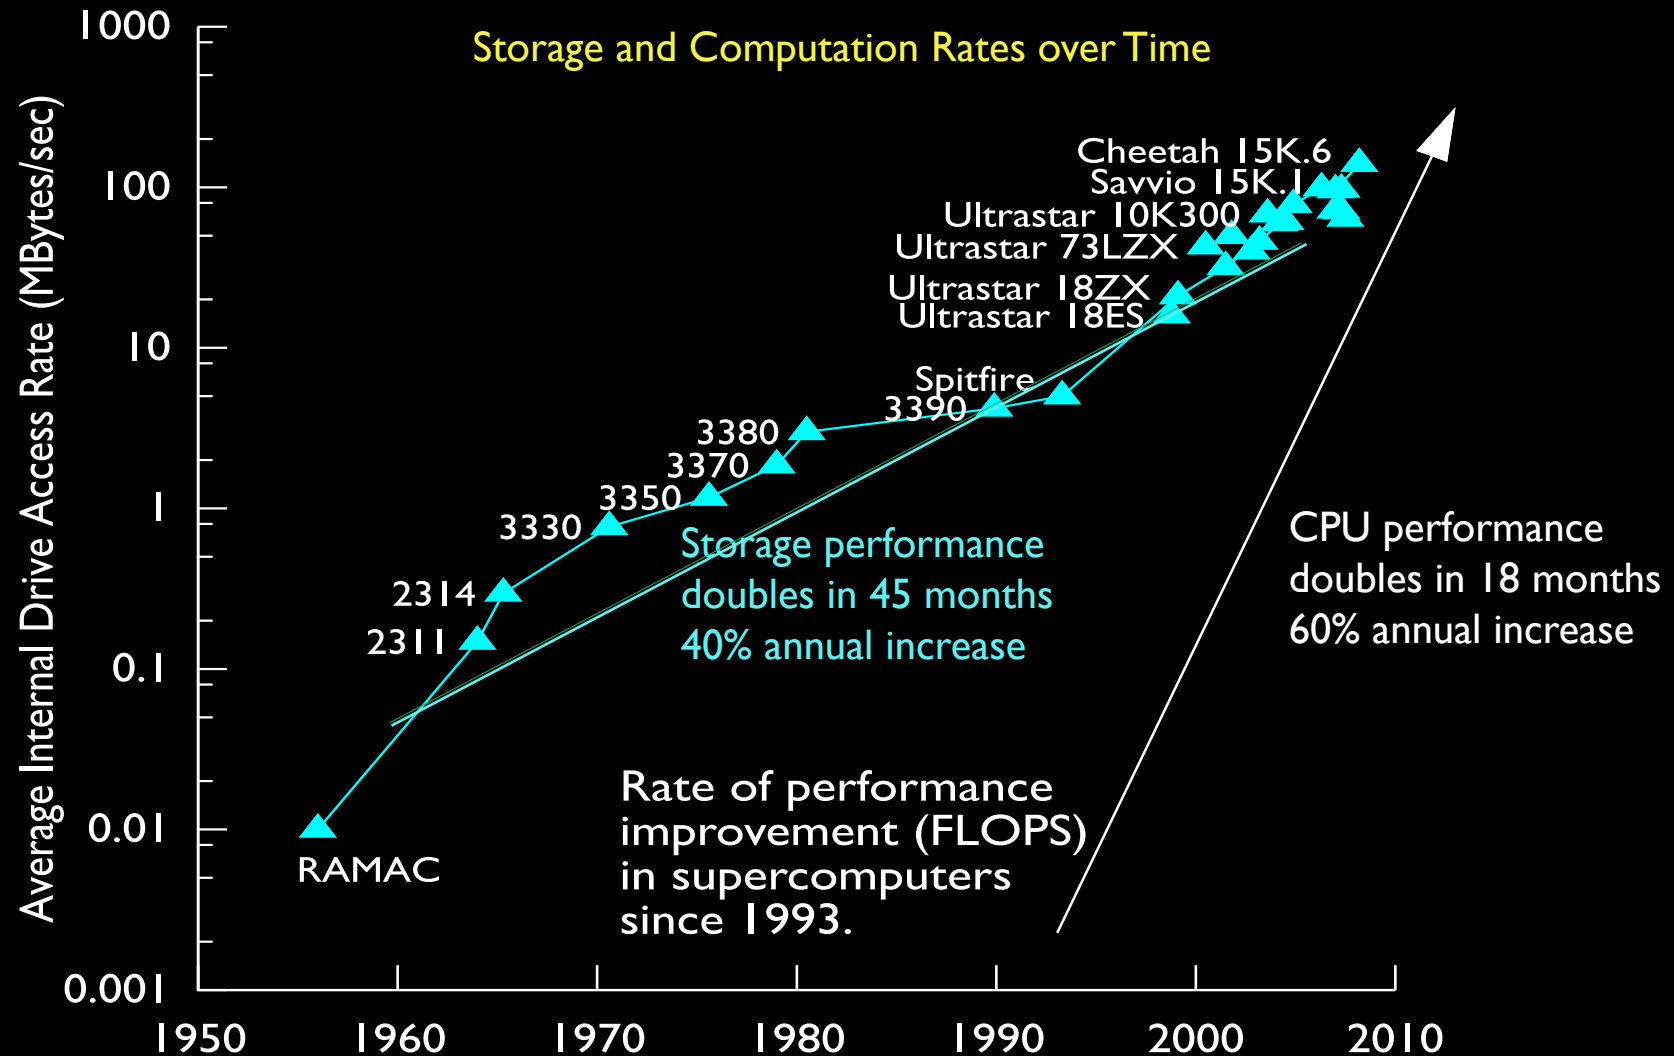
Tom Peterka

tpeterka@mcs.anl.gov

SIAM Minisymposium
February 26, 2010

Mathematics and Computer Science Division

We are computing more data, faster than we can manage.



Ref: Rob Ross, Visualization and Parallel I/O at Extreme Scale, SciDAC '08

More than Peak FLOPS: disk I/O rate limits analysis capability. Data that is not stored can't be analyzed.

Normalized Storage / Compute Metrics

Machine	Storage B/W (GB/s)	FLOPS (Pflop/s)	Flops per byte stored
LLNL BG/L	43	0.6	$O(10^4)$
Jaguar XT4	42	0.3	$O(10^4)$
Intrepid BG/P	50	0.6	$O(10^4)$
Roadrunner	50	1.0	$O(10^5)$
Jaguar XT5	42	1.4	$O(10^5)$

Percent Saved of Computed Data

Code	Domain	% Saved	PI
FLASH	Astrophysics	10	Ricker
Nek5000	CFD	1	Fischer
CCSM	Climate	1	Jacob
GCRM	Climate	10	Cram
S3D	Combustion	1-5	Bennett

Ref: CScADS Scientific Data Analysis & Visualization Workshop '09

-The average flops per byte of parallel I/O disk access today is between 10,000 and 100,000

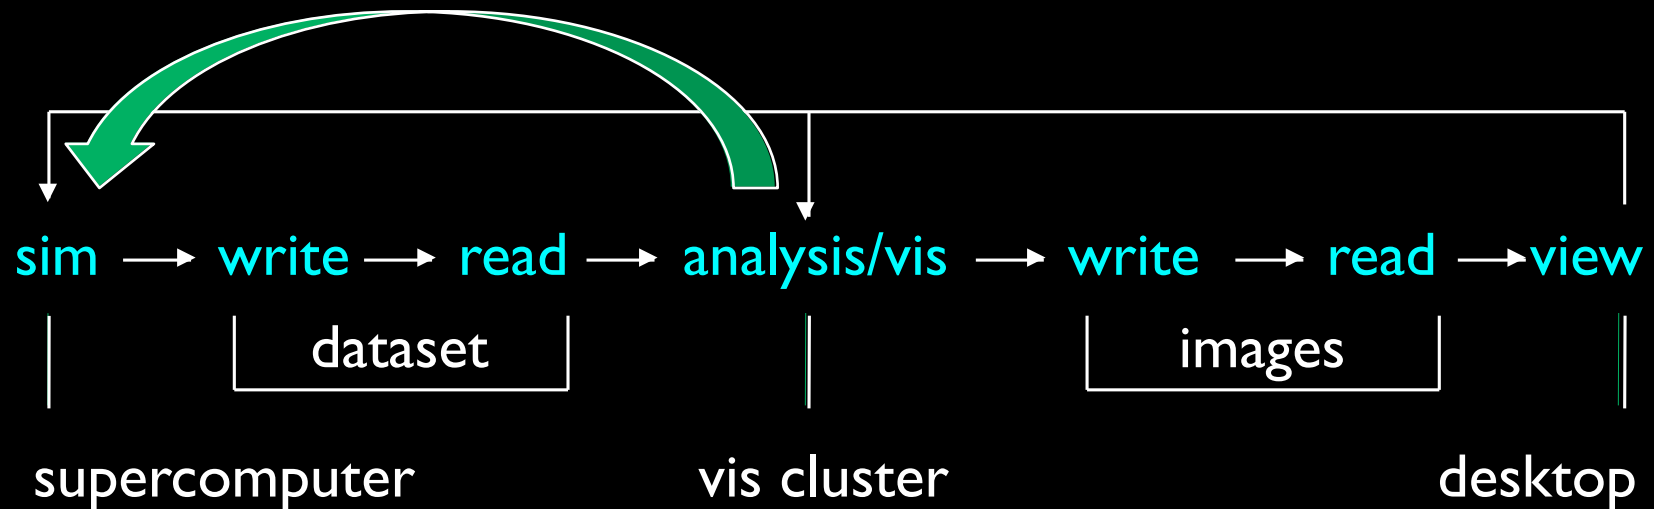
-In 2001, this number was approximately 500. Ref: John May, 2001.

-DOE science applications generate results at an average rate of 40 flops per byte of data. Ref: Murphy et al. ICS'05.

-Applications can only afford to save between 1-10% of what they compute.

-With postprocessing, what is not saved cannot be analyzed.

Our Science Workflow Cannot Scale Indefinitely

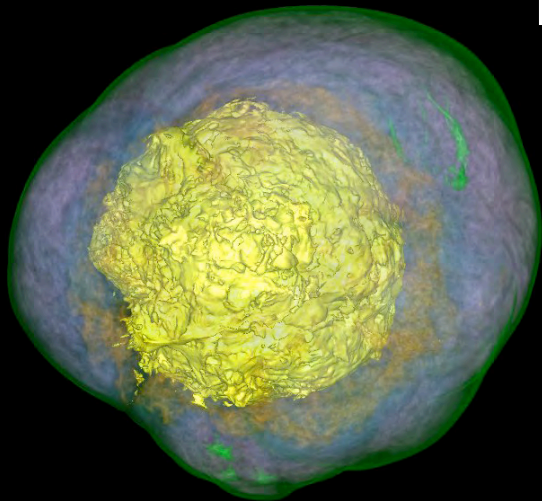


The increasing demands for analysis and visualization can be met by performing more analysis and **visualization tasks directly on supercomputers** traditionally reserved for simulation.

-Potential benefits: **Increased overall performance, reduced cost, tighter integration** of analysis and visualization in computational science.

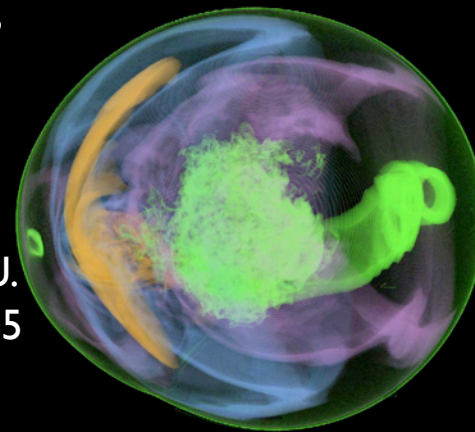
-Potential drawbacks: **Reduced per-core performance, increased load on computing resources, potential to crash computations.**

Parallel Volume Rendering

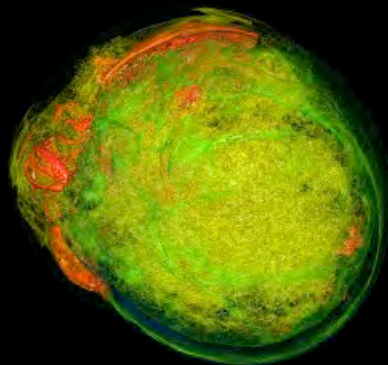


Pressure at time-step 1530

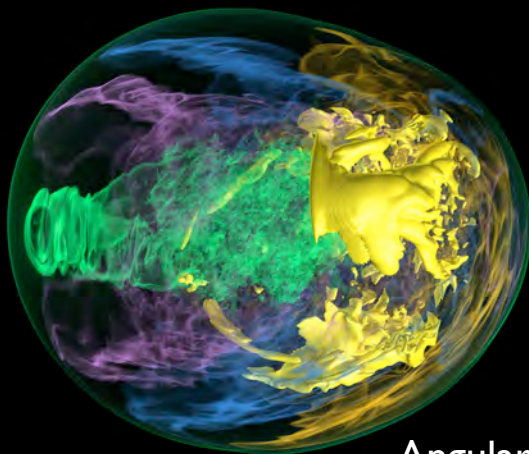
Volume rendering of shock wave formation in core-collapse supernova dataset, courtesy of John Blondin, NCSU. Structured grid of 1120^3 data elements, 5 variables per cell.



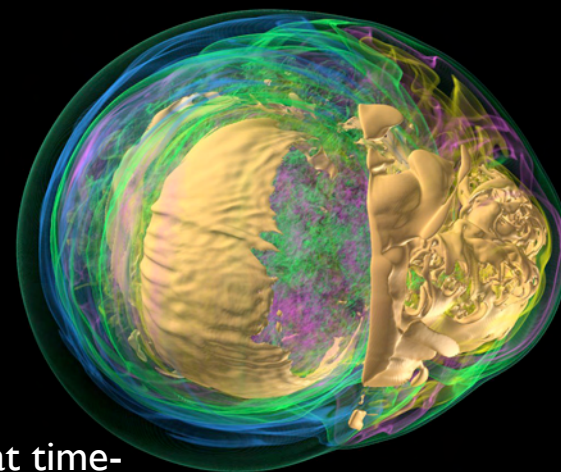
Angular momentum at time-step 1403



Entropy over 100 time-steps



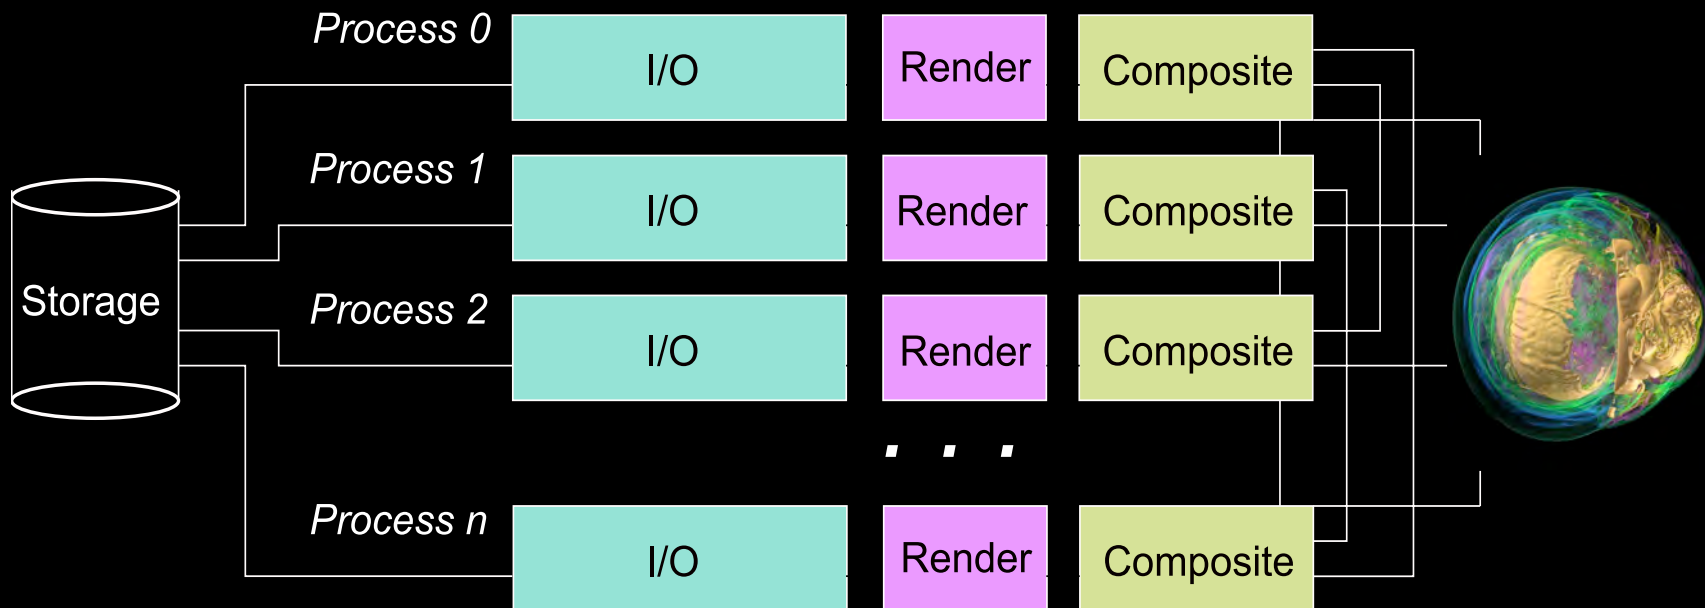
Angular momentum at time-step 1492



Entropy at time-step 1518

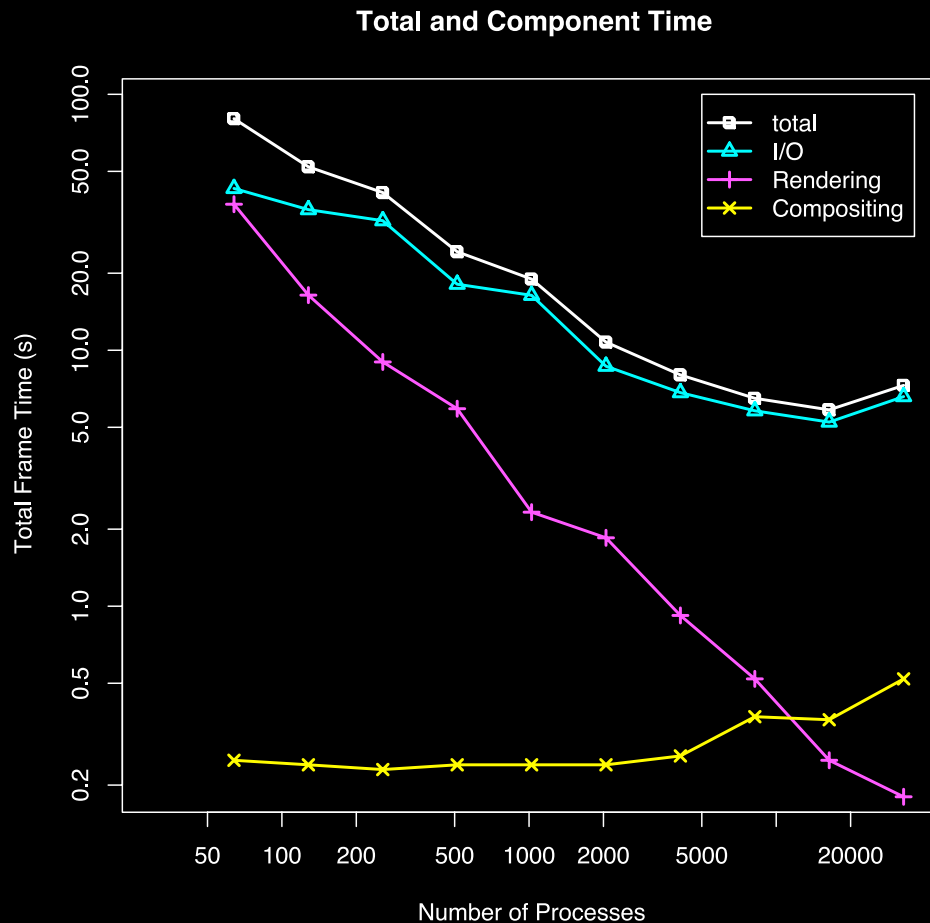
Parallel Volume Rendering Algorithm

Parallel structure for volume rendering algorithm consists of 3 stages performed in parallel

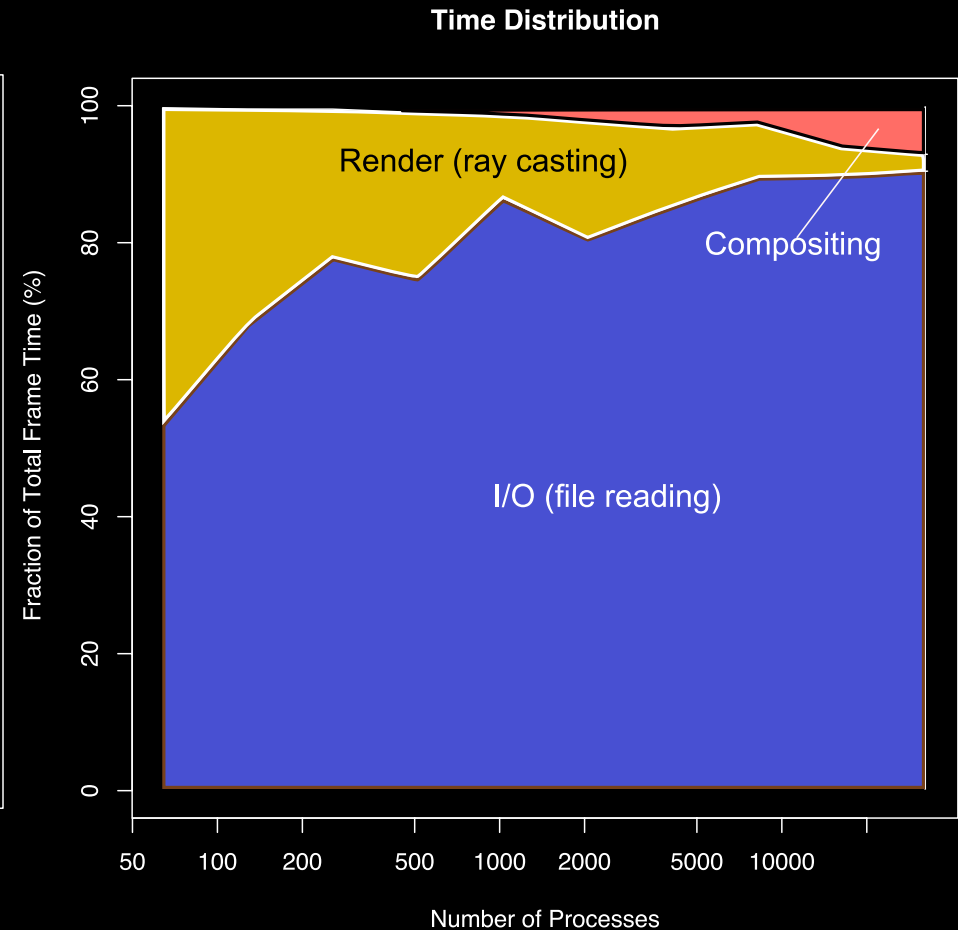


Parallel Volume Rendering on the IBM Blue Gene/P. EGPGV'08.

Performance: Total and Component Time



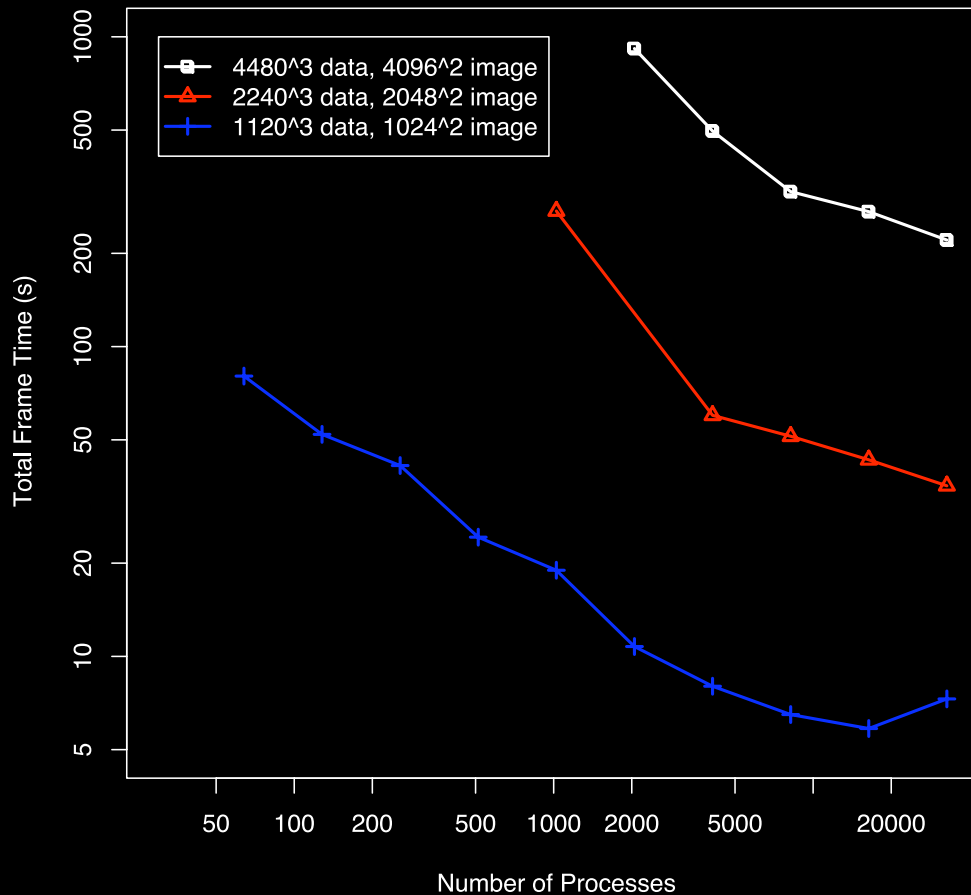
Total frame time and individual component times. Raw data format, 1120^3 , image size 1600^2 .



The relative percentage of time in the stages of volume rendering as a function of system size. Large visualization is primarily dominated by data movement: I/O and communication.

Performance: Large-scale Results

Volume Rendering End-to-End Performance



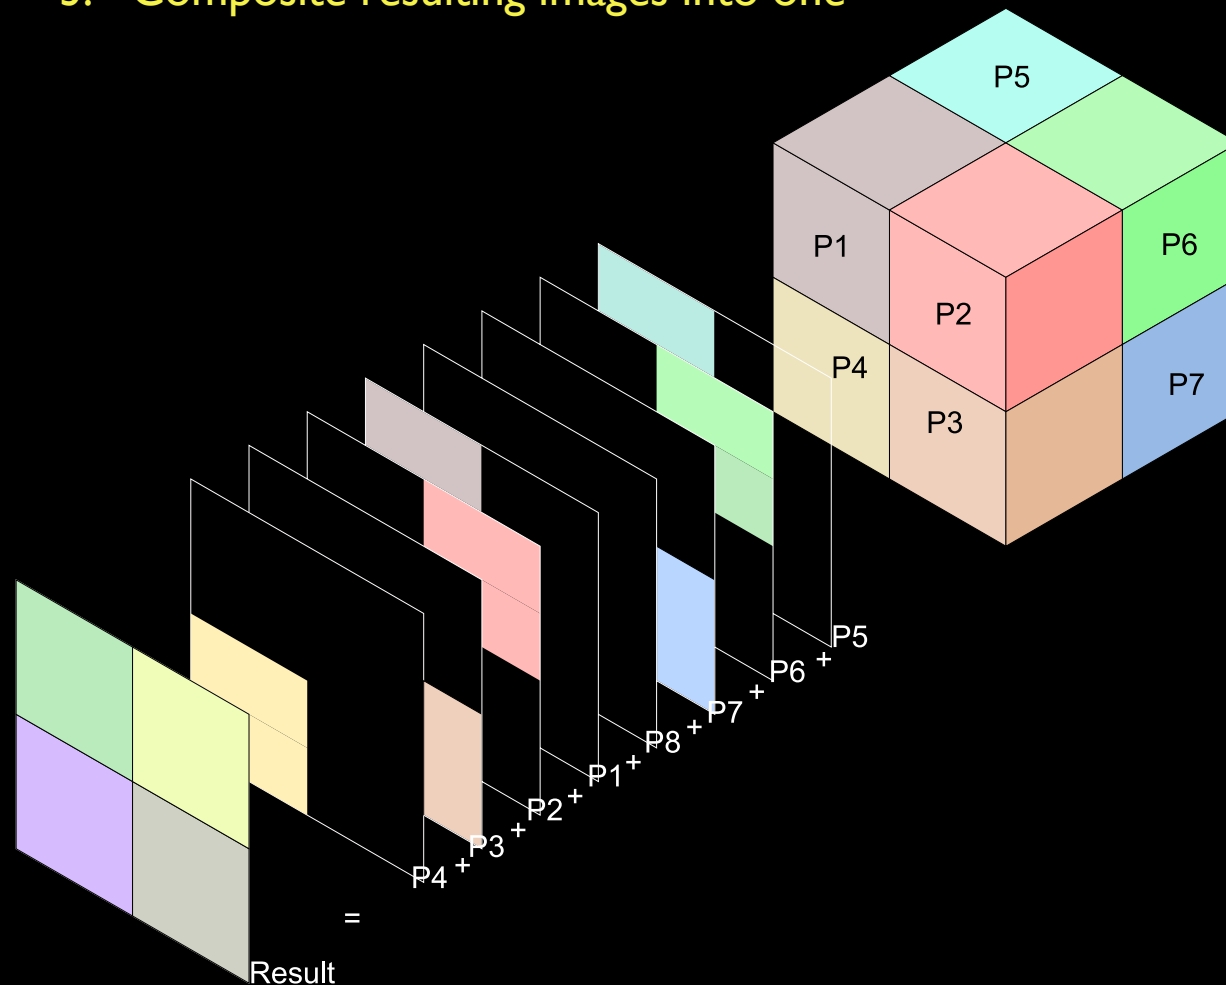
Grid Size	Time-step size (GB)	Image size (px)	# Procs	Tot. time (s)	% I/O	Read B/W (GB/s)
2240 ³	42	2048 ³	8K	51	96	0.9
			16K	43	97	1.0
			32K	35	96	1.3
4480 ³	335	4096 ³	8K	316	96	1.1
			16K	272	97	1.3
			32K	220	96	1.6

Scalability over a variety of data, image, and system sizes. A number of performance points exist for each data size.

Parallel Image Compositing

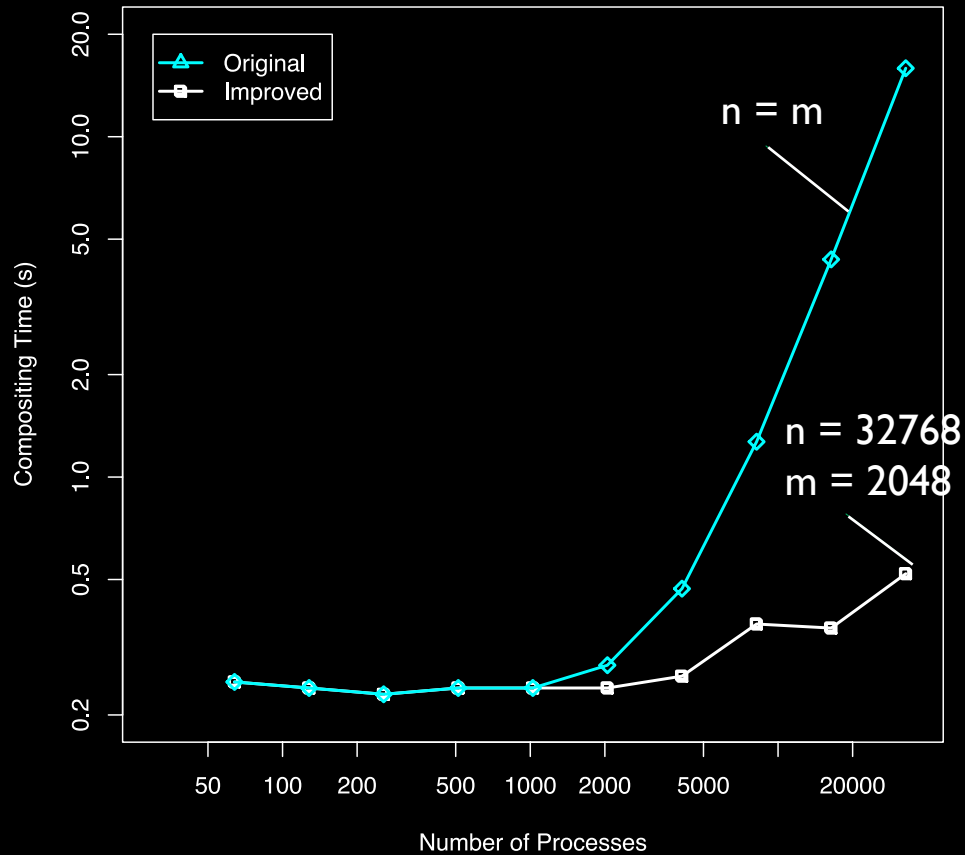
The final stage in sort-last parallel visualization algorithms:

1. Partition data among processes
2. Visualize local data
3. Composite resulting images into one

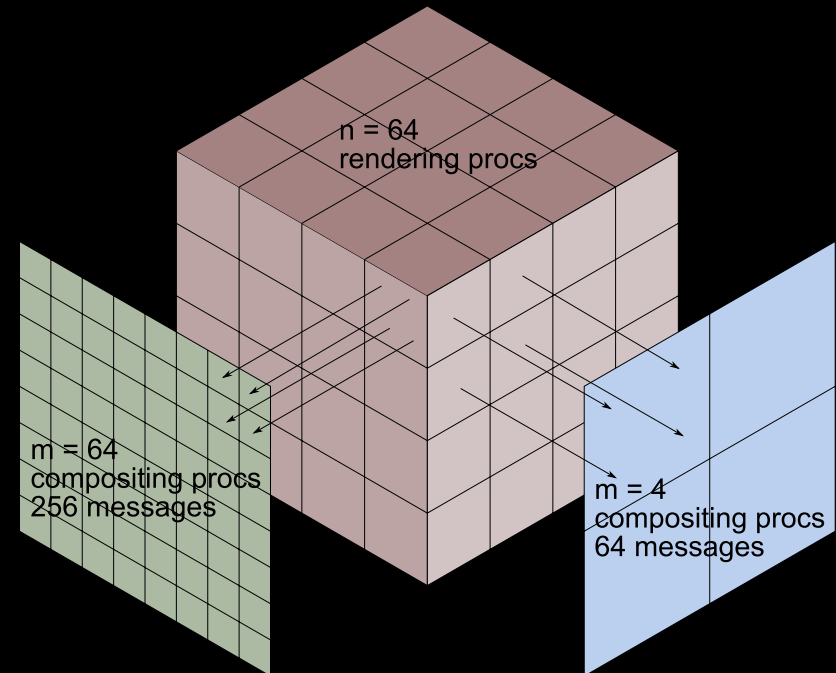


Direct-Send Optimization

Compositing Time



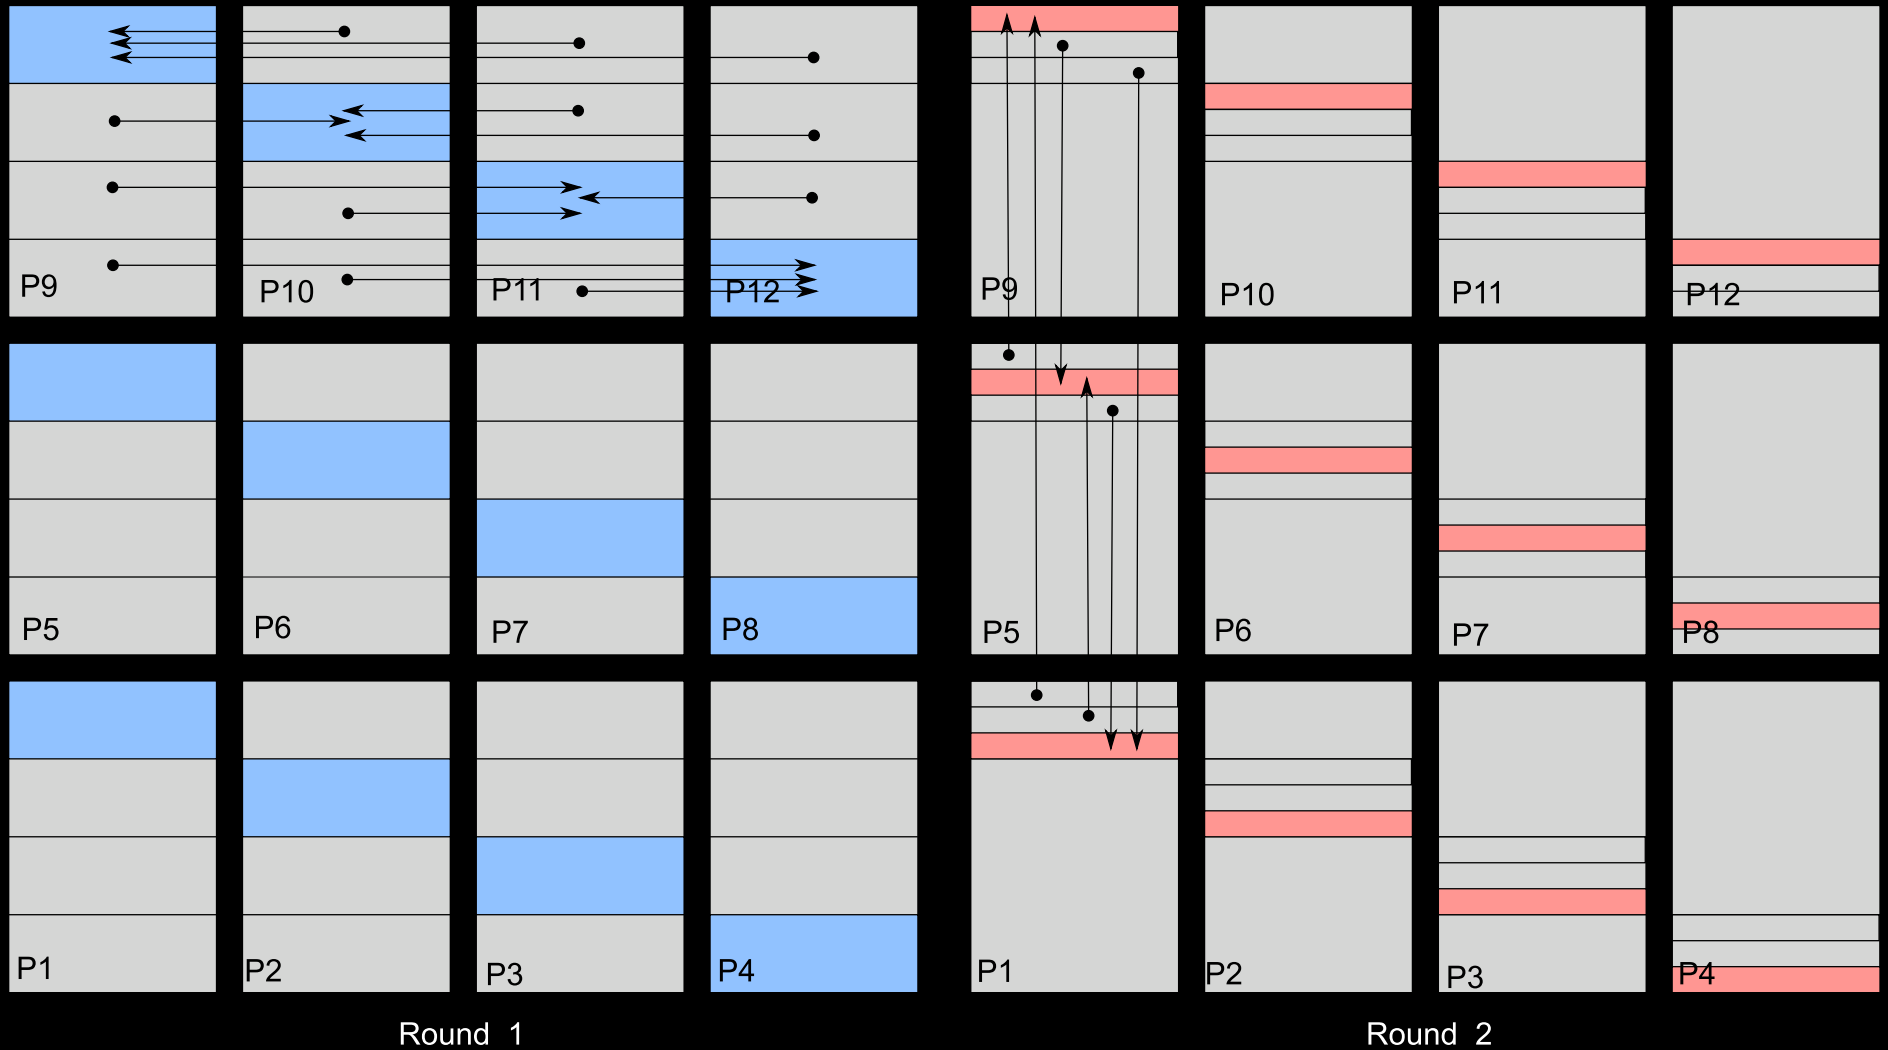
Direct-send compositing time improved up to 30X. | 120^3 data volume, 1600^2 image size.



Usually in direct-send, $n = m$, but setting $m < n$ can reduce contention when n is large. On average, $O(m * n^{1/3})$ total messages, can get down to $O(n)$ if $m = n^{2/3}$.

End-to-End Study of Parallel Volume Rendering on the IBM Blue Gene/P. ICPP'09

Radix-k Compositing Algorithm

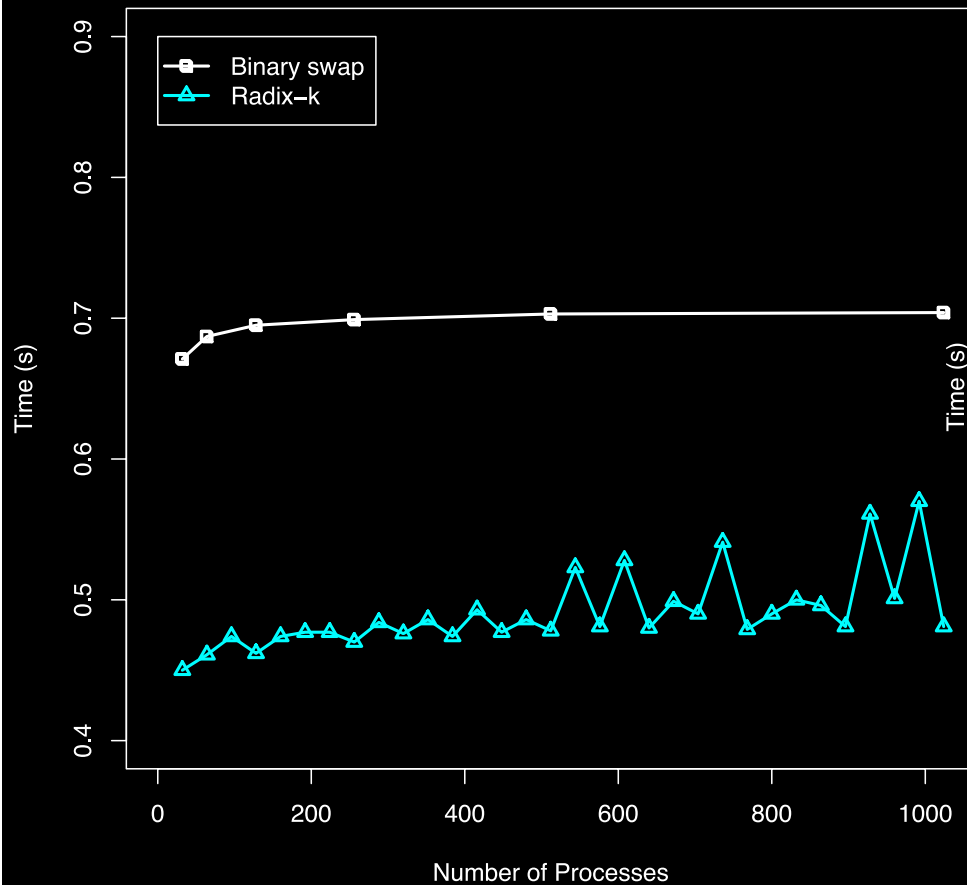


Radix-k: More parallel, managed contention, p does not need to be power of 2

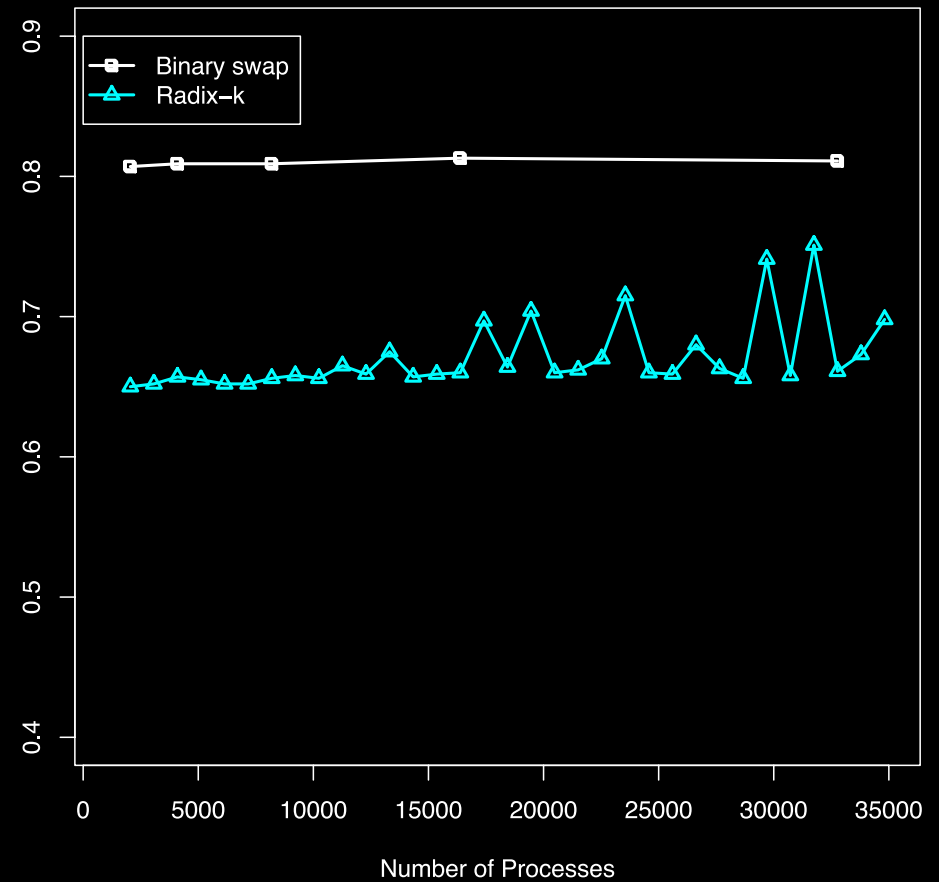
A Configurable Algorithm for Parallel Image-Compositing Applications. SC09

Radix-k Performance

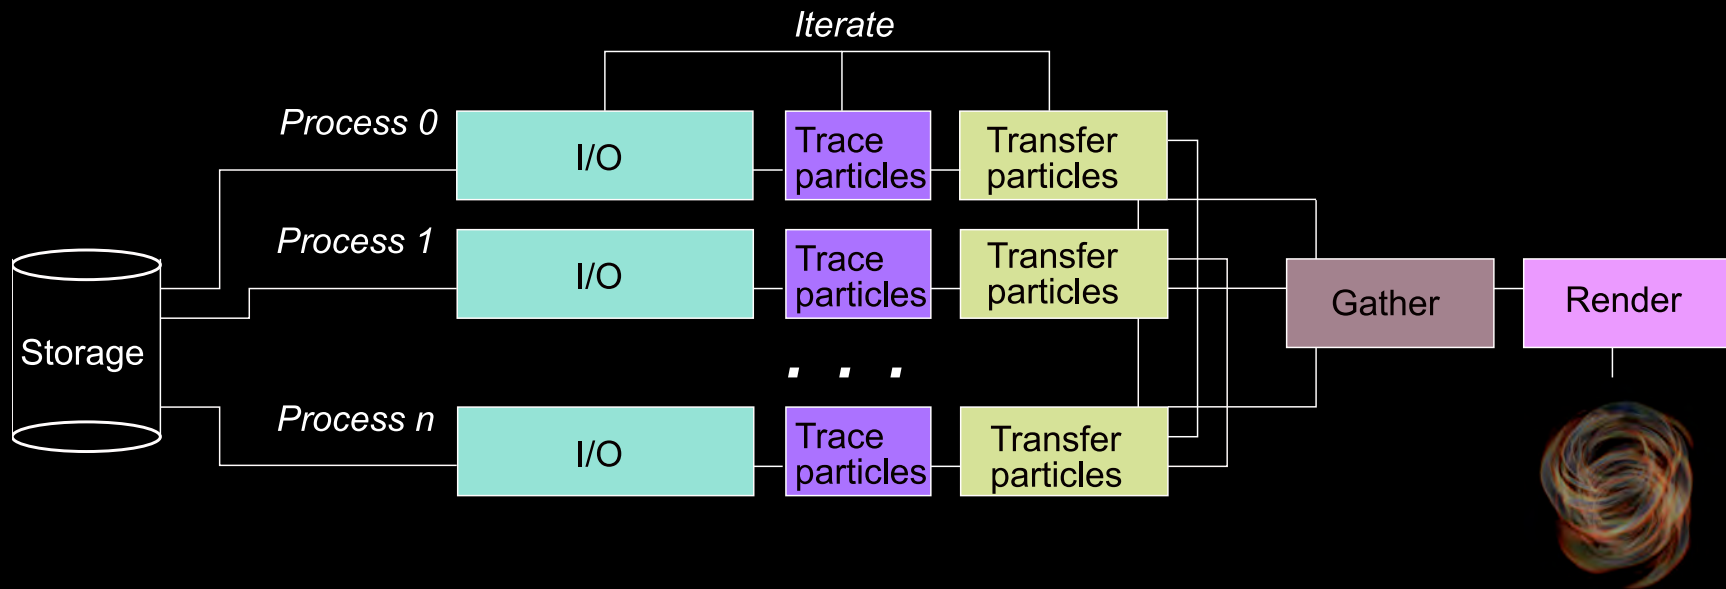
Compositing Time for 8 Mpx Image



Compositing Time for 8 Mpx Image



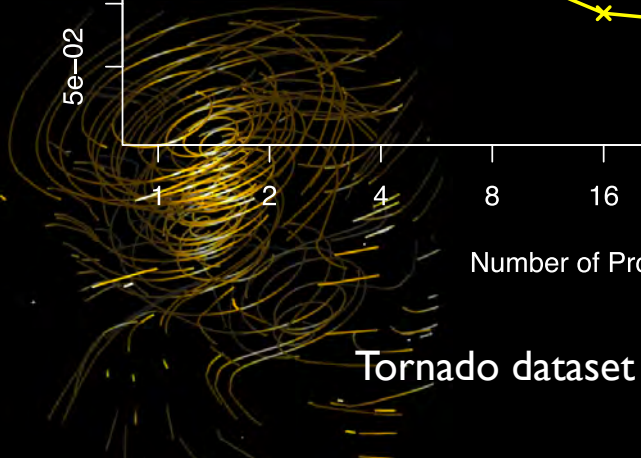
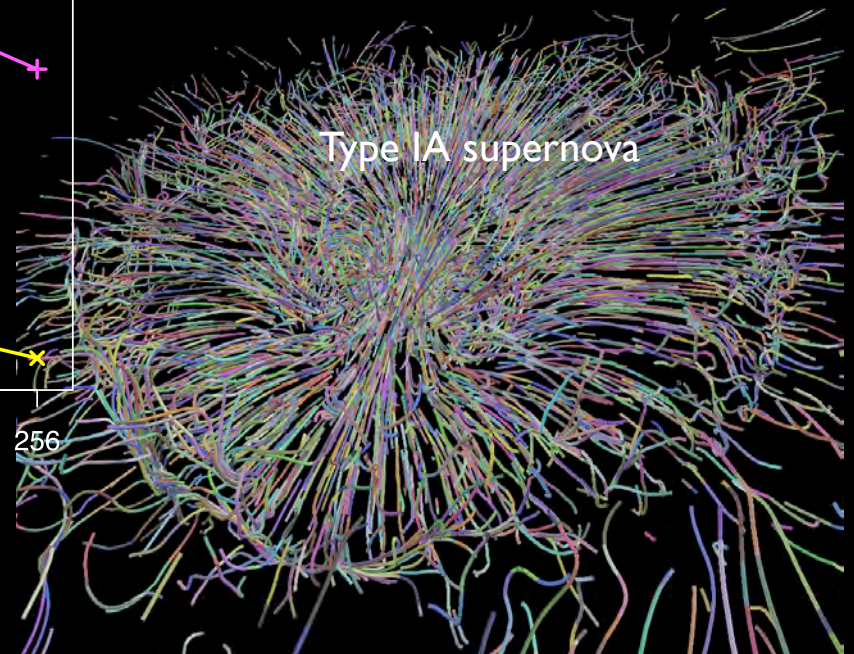
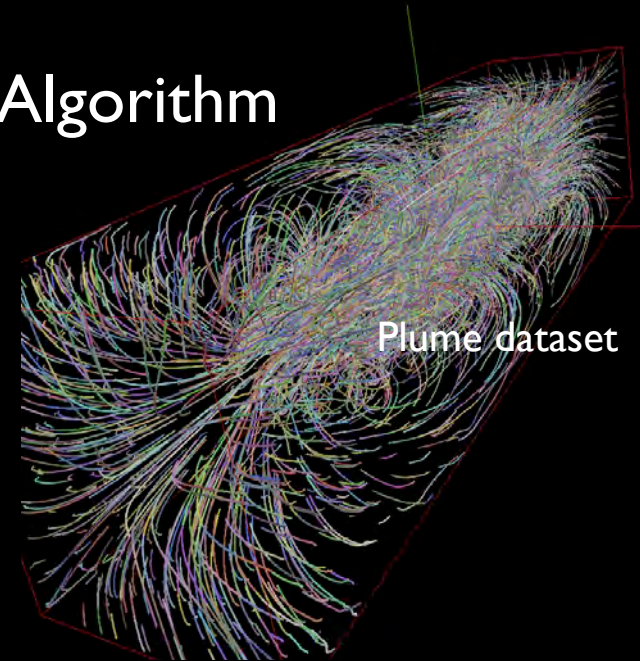
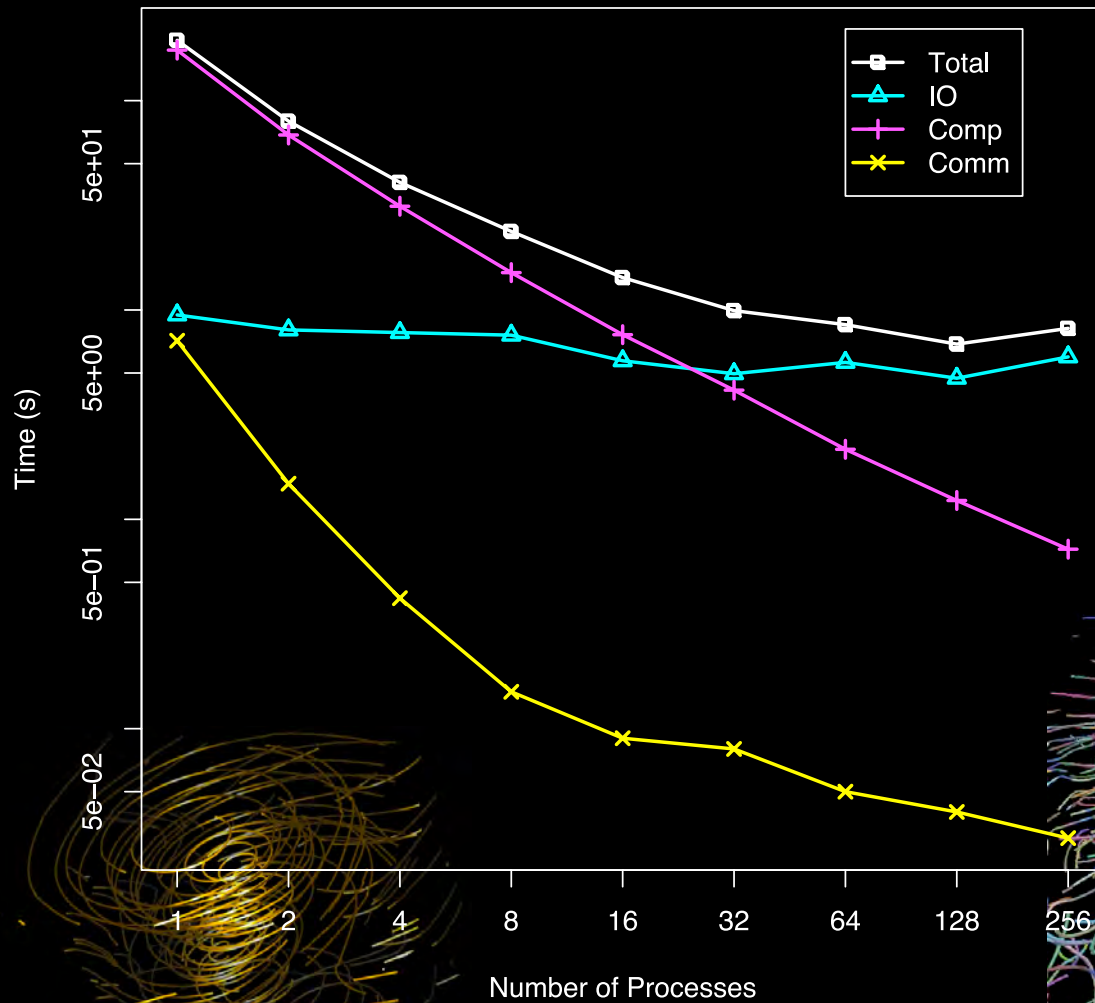
Parallel Flow Visualization Algorithm



Parallel structure for flow visualization algorithm consists of iterations of particle tracing and transfer, followed by a rendering stage.

Parallel Flow Visualization Algorithm

Time for 10 Rounds



Looking Toward In Situ Analysis & Visualization

Pros

- Reduced data movement
- Access to every data byte
- Native data structures
- Native algorithms
- Custom operations
- Increased accuracy

Cons

- Memory footprint
- Application constraints
- Increased complexity
- Expanded / collaborative domain knowledge

Challenges to Address

- Appropriate analysis / visualization applications
- Programming model
- Execution and use model

Further Reading

Peterka, T., Goodell, D., Ross, R., Shen, H.-W., Thakur, R.: A Configurable Algorithm for Parallel Image-Compositing Applications. Proceedings of SC09, Portland OR, November 2009.

Peterka, T., Yu, Hongfeng, Ross, R., Ma, K.-L., Latham, R.: End-to-End Study of Parallel Volume Rendering on the IBM Blue Gene/P. Proceedings of ICPP'09, Vienna, Austria, September 2009.

Peterka, T., Ross, R. B., Shen, H.-W., Ma, K.-L., Kendall, W., Yu, H.: Parallel Visualization on Leadership Computing Resources. Journal of Physics: Conference Series SciDAC 2009, June 2009.

Peterka, T., Ross, R., Yu, H., Ma, K.-L., and Girado, Javier: Autostereoscopic Display of Large-Scale Scientific Visualization. Proceedings of IS&T / SPIE SD&A XX Conference, San Jose CA, January 2009.

Peterka, T., Ross, R., Yu, H., Ma, K.-L.: Assessing Improvements to the Parallel Volume Rendering Pipeline at Large Scale. SC08 Ultrascale Visualization Workshop, Austin TX, November 2008.

Ross, R. B., Peterka, T., Shen, H.-W., Hong, Y., Ma, K.-L., Yu, H., Moreland, K.: Parallel I/O and Visualization at Extreme Scale. Journal of Physics: Conference Series SciDAC 2008, July 2008.

Peterka, T., Yu, H., Ross, R., Ma, K.-L.: Parallel Volume Rendering on the IBM Blue Gene/P. Proceedings of Eurographics Symposium on Parallel Graphics and Visualization 2008 (EGPGV'08) Crete, Greece, April 2008.



... for a brighter future



www.ultravis.org



U.S. Department
of Energy

UChicago ►
Argonne_{LLC}



A U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC

Moving Analysis to the Data: Scalable Visualization Using Simulation Resources

Acknowledgments:

Hongfeng Yu, Wes Kendall, Rob Latham, Dave Goodell, Kwan
-Liu Ma, Rob Ross, Han-Wei Shen, Rajeev Thakur
John Blondin, Tony Mezzacappa

Argonne and Oak Ridge Leadership Computing Facilities

US DOE SciDAC UltraVis Institute

Tom Peterka

tpeterka@mcs.anl.gov

SIAM Minisymposium
February 26, 2010

Mathematics and Computer Science Division